JOURNAL OF
Optical Communications
and Networking

PON going beyond FTTH [Invited Tutorial]

Thomas Pfeiffer,[1,*] **Pascal Dom,**[2] **Sarvesh Bidkar,**[1] **François Fredricx,**[2]
Konstantinos Christodoulopoulos,[1] **and René Bonk**[1]

[1] *Bell Labs, Nokia, Magirusstrasse 8, 70469 Stuttgart, Germany*
[2] *Fixed Networks Business Division, Nokia, Copernicuslaan 50, 2018 Antwerp, Belgium*
**Corresponding author: thomas.pfeiffer@nokia-bell-labs.com*

Received 30 July 2021; revised 29 September 2021; accepted 5 October 2021; published 2 November 2021 (Doc. ID 439241)

Time division multiplexing passive optical networks (TDM-PONs) are the most widely deployed optical system solutions in current broadband access networks worldwide. The energy and cost efficiency of both their implementation and operation has reached levels that also make them an attractive option for other cost sensitive communication networks. We discuss how essential features of current TDM-PON specifications can be leveraged to also use them for low latency and high capacity professional services in public and private networks. Also, possible PON architecture evolutions towards added intra-PON communication are outlined that are motivated by the latency requirements of some practical use cases.

https://doi.org/10.1364/JOCN.439241

1. INTRODUCTION

Time division multiplexing passive optical networks (TDM-PON) were originally conceived as a cost efficient solution for providing broadband services to residential users. The transmission protocol and the optical architecture were designed for supporting primarily best effort services and for exploiting statistical multiplexing gains. But business types of services with more stringent requirements on bandwidth (BW) and quality of service (QoS) were also taken into account from the beginning. Today, both the transmission convergence (TC) layer as well as the physical layer [physical medium dependent (PMD) layer] provide versatile and powerful means for configuring the transmission system to the needs of largely diverse services, and are continuously being further developed.

Despite their manifold capabilities, TDM-PONs have been deployed so far almost exclusively in public fiber-to-the-home or -building (FTTH/FTTB) networks for services to individual end users. With now more than 73 million optical line termination (OLT) ports and more than 900 million optical network unit (ONU) ports having been shipped worldwide since 2008 [1], the component and system costs have substantially decreased over time, despite increased system bit rates. As a result, TDM-PONs are currently explored also for use in other cost sensitive markets beyond FTTH that would benefit from the simplicity of a passive point-to-multipoint (ptmp) optical distribution network (ODN) architecture. One such opportunity is the market segment of business local area networks (LANs), which was addressed already several years ago. In passive optical LAN (POL), the distribution and aggregation of Ethernet data are accomplished by a TDM-PON for ptmp transport. The topology of office LANs is identical to PON, so that by replacing the Ethernet floor and workgroup switches with passive optical splitters and by dealing with active nodes only at the end-points of the network, the implementation and operation of the network are largely simplified. TDM-PON allows for a fully transparent transport of Ethernet traffic. The Ethernet data in LANs, as in any other use case, are encapsulated by the TC layer (layer 2) functions that ensure QoS aware transport over a ptmp architecture. Also, collision resolution, which in conventional LANs is accomplished by Ethernet switches, in this case is taken care of by the TC layer functions in the central OLT and in the distributed ONUs by organizing the optical transmission in time slots.

It is this principle of assigning BW on a predefined grid of time slots that offers TDM-PON to be used for deterministic, low latency, and low jitter data transport for non-residential services and use cases. We will consider in particular use cases in the following categories:

Wireless networks:
- back-/mid-/fronthaul transport in wireless radio access networks (RANs), especially for small cells in public and private areas.

Enterprise networks and vertical markets (with or without wireless):
- industrial production, manufacturing, and logistics.

Intra-datacenter (intra-DC) networks:
- small DC (not hyperscale), edge, and far edge cloud DC.

In each of those categories, there are uses cases that benefit not only from the ptmp architecture and efficiency of TDM-PON, but that would also benefit from direct communication between the end-points of the ODN (Fig. 1). Depending on the service category, the re-use of the ODN drop section for inter-ONU connections or for additional overlaid transport

Corrected 24 November 2021

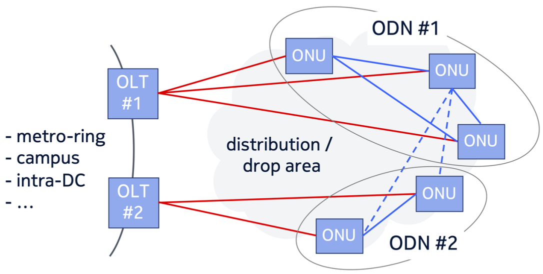

Fig. 1. Topology of intra-ODN (solid blue lines) and inter-ODN (dashed blue lines) communication in addition to the communication topology of conventional PONs (red lines). Some physical realizations of such local logical links on typical PON fiber infrastructures are proposed for different use cases in Section 5.

can help meeting low latency requirements and/or offload local traffic from the PON feeder section and from the OLT:

Wireless:

- fronthaul in a long reach PON,
- uplink joint reception in a multi-site multiple input multiple output (MIMO) scenario.

Enterprise:

- communication between sensors and actors within a manufacturing cell.

Intra-DC:

- east–west traffic between servers within a rack or within a pod.

Throughout this paper, we will consider the family of PONs specified by the International Telecommunication Union–Telecommunication Standardization Sector (ITU-T): single wavelength pair gigabit-capable PON (GPON) [2], 10-gigabit-capable symmetric PON (XGS-PON) [3], 50G-PON [4], and multi-wavelength channel next-generation PON 2 (NG-PON2) [5]. Also included in our considerations is 25GS-PON [6], which has not been specified by ITU-T, but reuses most of the TC layer specifications of XGS-PON. As far as the topic of this paper is concerned, those PON variants share similar TC layer specifications and apply common ODN definitions, but use mutually different wavelength ranges for transmission to allow for coexistence on the same ODN.

This paper is organized in the following way. In the next section, the general end-to-end network architecture will briefly be sketched, showing on a high level the interplay between applications and wireless and fixed network segments. Then those aspects of TDM-PON will be described that play an essential role in the latency and jitter performance of the system. Next will be the discussion of the interplay between service and TDM-PON resource assignments and possible means for improvement. Finally, enhancements of the ODN and (in some cases) of the system optics will be introduced that will allow for efficient intra-ODN and inter-ODN communication.

2. APPLICATION AND TRANSPORT NETWORK ARCHITECTURE

The generic network architecture shown in Fig. 2 comprises the essential elements and segments that contribute to the end-to-end latency on the application layer (T_{app}), assuming that the application is provided over a wireless network with fixed network segments for x-haul transport. On the wireless side, there are the 5G core (5GC), here showing only the user plane function (UPF) of it, and the functional blocks of a split base station architecture: central unit (CU), distributed unit (DU), and radio unit (RU), which may all or partly be virtualized on a compute infrastructure, or be implemented in dedicated hardware. They can be residing in different or pairwise common locations or all combined in one location. The backhaul, midhaul, and fronthaul transport links (x-haul in short) are assumed to be based on TDM-PON technology for the discussion in this paper. The links add delays for backhaul (T_{BH}), midhaul (T_{MH}), and fronthaul (T_{FH}), respectively, due to fiber propagation and due to PON protocol and system implementation. The wireless propagation over the air adds the delay T_{air}. The application devices (here server and client) may be attached to the UPF and to the user equipment (UE) directly or via interworking functions (or gateways) on either side, which ensures compatibility of data and transport protocols between both layers. They introduce additional latencies at the server side (T_{IWFs}) and at the client side (T_{IWFc}), which for this discussion comprise both processing delays and propagation delays induced by external cables. All the listed latencies can change over time, depending on the individual root cause, e.g., buffer delays (PON protocol or store-and-forward induced) can vary with packet size on a lower microsecond scale, whereas fiber propagation delays vary with temperature on a lower nanosecond scale or even less for the link lengths considered here. The transport delays per x-haul segment will in general depend on the location of the ONU within the ODN as well as on the resource assignment details for that particular ONU. So the parameter values for the architecture in Fig. 2 will change on a per-ONU basis, and over time.

The computing and processing delays in the 5GC, CU, DU, RU, and UE are not considered in detail for this transport related discussion.

In general, backhaul and midhaul latencies are considered to be in the range of single or double digit milliseconds, whereas fronthaul latencies are in the sub-millisecond range (typically

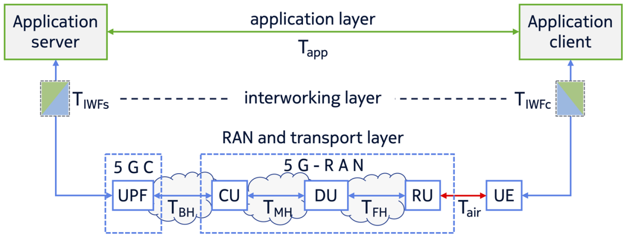

Fig. 2. Generic network architecture for an end-to-end application served over a 5G network, containing fixed network transport segments between functional blocks of the split base station architecture (for the latency related time labels, see text).

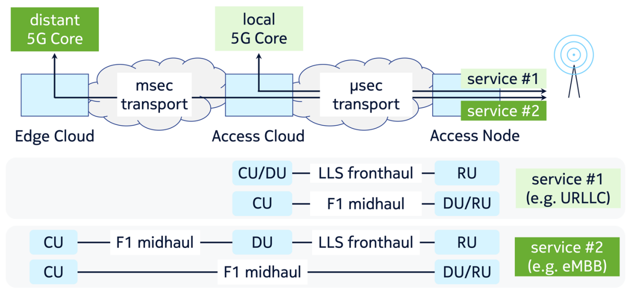

Fig. 3. Combined network architecture containing transport links and RAN processing functions that can support low latency and latency tolerant services on a common infrastructure (LLS, low layer split as defined by O-RAN [7]; F1, high layer split as defined by 3GPP [8]). If a service requires an end-to-end latency of only one or a few milliseconds, then any transport latency would typically be constrained to the sub-millisecond range (cf. Fig. 2). In the above architecture, that even applies to an F1 midhaul link for service #1.

100–150 μs one way). The latter requirement has been introduced by the hybrid repeat request (HARQ) process that is needed for correcting erroneous transmission of data across the air interface. The details vary between LTE and 5G networks; however, the sub-millisecond fronthaul requirement remains valid in both cases, whenever HARQ is enabled.

The millisecond range for backhaul and midhaul is acceptable for enhanced mobile broadband (eMBB) services, but may not be acceptable, when certain low latency services are considered, such as ultra-reliable low latency communication (URLLC) services or machine-to-machine (M2M) type communication services. In such cases, it may be necessary to terminate the user traffic close to the application, thus keeping T_{BH} and/or T_{MH} sufficiently small. Other services running on the same network, or other less critical parts of the URLLC service, may still allow for higher latency values. A combined network satisfying both requirements while using a common air interface is shown in Fig. 3. Here the low latency service #1 (e.g., M2M control loop) is terminated via the local 5GC in the Access Cloud, while the latency tolerant service #2 (e.g., video analytics for general process monitoring in a factory) can be terminated in a more powerful server via a distant 5GC in the Edge Cloud.

The access node (i.e., the radio and antenna site in Fig. 3), Access Cloud, and Edge Cloud may host the functional blocks of the 5G system in various ways as indicated in the figure. The transport segments between them are laid out for meeting the latency requirements, and hence are defined by their respective transport delays rather than by specific distance values.

Other mixed architectures are possible as well, as will be shown in the last section of this paper.

3. LATENCY COMPONENTS IN TDM-PON

For discussing latencies in transport over TDM-PON, the following processes must be considered, with their impact briefly summarized below. More details for the first three categories are provided in the sub-sections thereafter, following the process descriptions in [2,5].

Dynamic BW assignment (DBA) updates:

- zero delay for fixed BW assignment per ONU,
- close to zero delay for pro-active assignment [e.g., cooperative DBA (Co-DBA), or pre-configured industrial Internet of Things (IIoT); see next section],
- up to 1 ms delay for re-active assignment (e.g., best effort services).

ONU buffering for accommodating the scheduled start time of burst transmission:

- max 7.8 μs delay (with 16 periodic bursts per frame) up to, e.g., max 1 ms delay (with one burst every eight frames).

ONU activation and ranging introduces quiet windows:

- zero delay for pre-configured networks (IIoT), or with second optical channel dedicated to ranging,
- <10 μs delay, if ONUs are pre-registered and their location is approximately known to within 1 km differential distance,
- up to a few 100 μs delay for general (current) use cases.

PON equipment and fiber link:

- PMD layer digital processing delays:
  - equalization for dispersion and equipment BW limitations: <1 μs for maximum likelihood sequence estimation (MLSE),
  - (de)scrambling: a few ns,
  - forward error correction (FEC) : <5 μs end-to-end for low density parity check (LDPC) code.
- Ethernet switching and data processing delays:
  - store-and-forward: 1.2/7.2 μs (0.5/2.9 μs) for 1500/9000 byte Ethernet packets at 10 Gb/s (25 Gb/s),
  - queuing delays: depending on total load, service types, and priority settings.
- Fiber propagation:
  - 5 μs/km delay approximately.

A. Dynamic Bandwidth Assignment

The dynamic bandwidth assignment (DBA) process in TDM-PON is conceptually composed of two main components: the BW assignment to each service (transmission container, T-CONT) per ONU, and the BW map (BWmap) generation for organizing the upstream (US) data transport for all ONUs and services in an US frame according to the BW assignment (Fig. 4).

The BW assignment per service takes into account parameters that are defined when setting up the service, as well as PON-internal adjustment parameters.

The service related parameters (specified only with service setup) for the per service BW assignment are

- R_F, R_A, and R_M for capacity (fixed, assured, maximal BW);
- χ_{AB}, P, and ω for priority (indication of priority in the case of congestion);
- T_{JT}, T_{DBT}, and T_{PST} for timing (jitter tolerance, BW assignment delay tolerance, protection switching delay tolerance).

PON internal BW adjustments (re-active, dynamically updated) are derived from

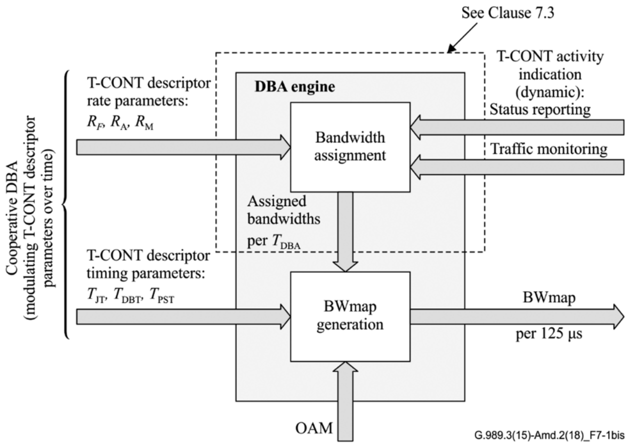

G.989.3(15)-Amd.2(18)_F7-1bis

Fig. 4. Schematic overview of the DBA process building blocks in TDM-PON (Fig. 7-1bis in [5]).

- reports of ONU buffer filling status (status reporting),
- utilization of previously assigned capacity (traffic monitoring).

A new assignment process has only recently been defined: cooperative DBA (Co-DBA), e.g., accommodating transport of dynamically changing fronthaul traffic in RAN. It proactively assigns future US BW based on dynamic requests from the wireless DU function. It will be described in Section 4.B.

The dynamics of the BWmap generation are dependent on the respective service and on the hardware and/or software implementation of the DBA process:

- it can be static for fixed service BW, or
- it can be changed with every Nth frame (e.g., N = 1 or 8 or . . .), i.e., it gets updated every 125 μs or every 1 ms, etc.

The BWmap, whether changed or not, is sent to all ONUs with every downstream (DS) frame.

An important feature for latency sensitive services is the option to assign to each ONU up to 16 bursts per 125 μs frame, which allows for serving an ONU every 7.8 μs. On the other extreme, e.g., one burst can be assigned to each ONU only every eighth frame, i.e., once per millisecond, or even less frequently. The optimal assignment frequency is chosen depending on the trade-off between ONU buffer latency and US transmission BW efficiency. It can be different for different ONUs and for different services in the same PON, and in general does not need to be periodic.

For conventional best effort FTTH or POL services, the BW assignment per service (per T-CONT) is typically updated on a single digit millisecond time scale (DBA cycle time T_{DBA}), based on the parameters and internal updates as described above. For ultra-low-latency services, however, this will be too slow. But also for such cases, the service BW assignment and BWmap generation can still be modified within the above framework for accommodating those latency requirements. This will be shown for some examples in Section 4.

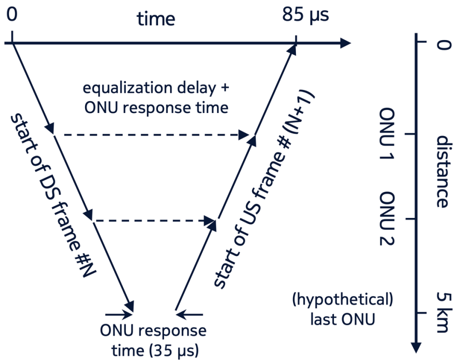

Fig. 5. Propagation of DS and US frames in a short (5 km) TDM-PON.

B. Physical Layer Frames and Bursts

The DS and US transmission in ITU-specified TDM-PONs is organized in frames of 125 μs duration. While the DS frames are real frames, the US frames are only virtual entities for providing a common reference for the relative timing of all US bursts. For the PON flavors primarily considered in the following (XGS-PON and 25GS-PON), those US frames are divided into 9720 time slots of 12.86 ns duration. The actual BWmap for US transmission is sent to all ONUs with every DS frame. It contains the assigned starting time slots and grant sizes of all bursts in the next (virtual) US frame. The US frame #(N + 1) starts propagating towards the OLT 35 μs (= fixed response time of all ONUs) after the DS frame #N has arrived at the farthest (real or hypothetical) ONU in the system. The longest possible distance between the OLT and ONUs in a given PON must be fixed when setting up the system. In most current FTTH deployments, it is typically set to 20 km or 40 km. In that case, the US frame #(N + 1) will arrive at the OLT 235 μs or 435 μs, respectively, after the DS frame #N has left the OLT. This implies that the ONUs will be able to apply a new burst timing and BW in the worst case (ONU close to OLT) only almost 235 μs after the BWmap has been updated by the OLT. However, when low latency services, including fronthaul, are considered, then the maximal OLT–ONU distance can be reduced to, e.g., 5 km or less, depending on the acceptable fiber propagation delay or on the extent of the network environment (e.g., campus networks). With 5 km maximal distance, the burst timing update would be implemented by the ONUs in the worst case up to 85 μs after the updated BWmap has been broadcast to the ONUs (Fig. 5). However, to be precise, this update delay does not mean that a user data packet would have to wait that long when it arrives at the ONU. For low latency traffic, one or multiple bursts are assigned to the ONUs in every US frame, so that there is always a timely transmission opportunity with an assigned base BW. Only potential modifications of the BW assignment would be affected by the above update delay. (For simplicity, we have assumed the ONU response time to be 35 μs for all ONUs in the above discussion. The standards, however, allow for a ±1 μs variation across all ONUs in the network, as long as they are constant over time, e.g., [3]).

C. ONU Activation and Ranging

For each ONU, the virtual US frame #(N + 1) starts at a time equal to "equalization delay + ONU response time" after receiving the DS frame #N. Hence, the correct setting of the equalization delay per ONU is an essential precondition for collision-free transmission of bursts in each US frame. It is determined and assigned by the OLT during the activation and ranging process for each ONU. In conventional FTTH networks, this process involves interrupting the US traffic twice, each time for up to 250 µs in a 20 km PON. These long interrupts are detrimental for low latency services. They can be shortened to only a few µs, if the fiber distance to new ONUs joining the system is approximately known *a priori* to within, e.g., 500 m accuracy, either by using ODN deployment data or by employing an alternative activation scheme as described in [9]. Another option is using a second PON channel for this process, either within a group of NG-PON2 channels or by another PON variant operating in a different wavelength range.

The equalization delay is determined and assigned with a granularity of 402 ps for XGS-PON and 25GS-PON. It can be corrected as needed during system operation, whenever the OLT receiver detects a drift of the US bit arrival times exceeding a configurable threshold value.

The initial determination of the equalization delay is based on a round trip time measurement per ONU during the ranging process. The propagation delay asymmetry between DS and US direction is taken into account by introducing the ratio of group refractive indices. For the basic wavelength set of XGS-PON (DS: 1577 nm, US: 1270 nm) the DS transmission takes longer than the US transmission by ~1.6 ns per fiber kilometer.

D. Time Alignment in PON

Aside from avoiding burst collisions, the equalization delay is also needed for time of day (ToD) synchronization of the ONU clocks to the OLT clock. The OLT sends a reference timestamp to the ONUs via the ONU management and control channel (OMCC), along with the DS frame number associated with it. The ONUs can then derive the respective ToD value from this timestamp by subtracting their precise ONU response time and equalization delay, again corrected for DS versus US link asymmetry and for one way instead of round trip delay. The inaccuracy of the derived ToD value is typically less than 10 ns.

The above ToD precision applies only to the PON internal protocol synchronization. When talking about total ToD error on the ONU side, the integration of the PON into an external network must also be considered. The clock frequency of the OLT chassis is typically synchronized via synchronous Ethernet (SyncE) to an external reference frequency derived from the incoming Ethernet data stream. Likewise, the phase and ToD for the chassis clock is synchronized to an external source via the IEEE 1588v2 Precision Time Protocol (PTP). The OLT port takes these references and forwards them to the ONUs: for ToD via OMCC as described above, for clock frequency via the 8 kHz DS frame rate. On the ONU side, a new SyncE and PTP based synchronization path is set up towards the client system. Taking all the involved processes in the nodes into account, the end-to-end performance of the PON link is capable of meeting the accuracy requirement of a class B telecom boundary clock (T-BC) with constant time error |cTE| < ±40 ns, consisting of two media converters (Fig. 6).

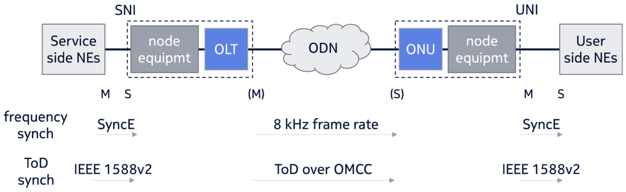

Fig. 6. Frequency and ToD synchronization across TDM-PON employing SyncE and IEEE 1588v2 outside of the PON and using the DS frame rate and OMCC inside the PON.

4. LOW LATENCY TRANSPORT FOR FRONTHAUL AND IIoT SERVICES

In this section, it will be shown how TDM-PON can be configured to support low latency transport for fronthaul in wireless RAN as well as low latency applications in IIoT networks.

A. Constant Bit Rate and Strictly Periodic Service Traffic

Fronthaul transport at the option 8 split point in 4G centralized RAN (C-RAN) employing the Common Public Radio Interface (CPRI) protocol is an example of constant bit rate traffic with sub-millisecond latency requirements. In the proof of concept described in [10], commercial equipment was used both on the wireless side (4G LTE radio head as a small cell) and on the PON side (XGS-PON). Bidirectional CPRI data at 2.5 Gb/s constant bit rate were transported over the PON that was configured for fixed BW, using periodic bursts in US. With a single burst per 125 µs frame per ONU, a round trip time of up to only 200 µs as required by the wireless equipment could hardly be accomplished, leaving no margin for fiber transport. However, with four bursts per 125 µs frame per ONU, up to 6 km of fiber could be bridged as enabled by the shorter buffer time at the ONU (one burst per ONU every 31.25 µs). The measured round trip time included processing in additional gateways that were used for encapsulating the CRPI data into Ethernet frames (and reverse) to accommodate the input format for the PON system (pre-standard radio-over-Ethernet as specified later in IEEE 1914.3). The encapsulation of the continuous CPRI stream at the input gateway generated periodic Ethernet frames of alternating sizes. At the output gateway, the continuous CPRI stream was fully recovered with sufficiently small jitter for proper operation of the small cell.

A similar case of deterministic periodic data transmission is encountered in the industrial space. M2M communication on the manufacturing floor is currently done using fieldbuses or industrial Ethernet protocols. In many current practical implementations, machine sensors and actors exchange small data packets at regular intervals. As an example, small groups of up to 50 nodes exchange data packets of up to 64 bytes every 30 ms or even faster than every 1 ms, at line rates of 100 Mb/s

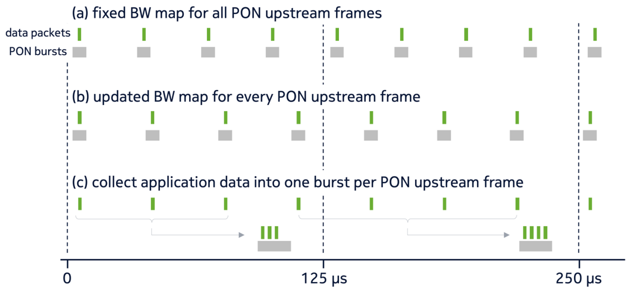

Fig. 7. Upstream transport of periodic application data, where the application data period (a) matches a sub-multiple of the PON frame period or (b), (c) does not match. In (b), the burst period is set equal to the data period with BWmap updated with each PON frame. In (c), multiple application data packets are aggregated in a single burst per PON frame. The data period is recovered by gateway buffers.

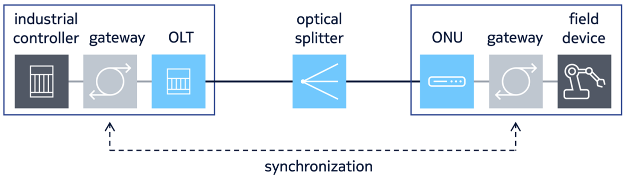

Fig. 8. Fieldbus or industrial Ethernet communication across TDM-PON, using gateways for encapsulation and possible aggregation/disaggregation of client data packets, as well as for jitter compensation.

or 1 Gb/s [11]. For the most challenging isochronous use cases, the required latency is less than 1 ms and the jitter less than 1 µs. The TDM-PON transport scheduling can be matched to such services. In some cases, the periodic transmission pattern of the client devices can be directly mapped to a matching periodic burst pattern on the PON in each frame Fig. 7(a)]. More likely, however, is the case that the client data, although periodic, would not match a fixed burst pattern that is repeated in every PON frame. Instead, the periodic client data would need a correspondingly periodic burst pattern that, however, needs a BWmap updated for each PON frame [Fig. 7(b)]. As a third option, multiple client data packets would be collected into one burst per frame and would be distributed again after transmission over the PON, thus recovering the original strictly periodic data stream [Fig. 7(c)]. As in the case of CPRI transport over TDM-PON, this adaptation would be accomplished using appropriate gateways or interworking functions (Fig. 8).

B. Dynamically Changing Service Traffic

Fronthaul transport at the option 8 interface, using the CPRI protocol, was identified early on to poorly scale with network capacity and antenna configurations in 5G networks. With an intra-physical layer (intra-PHY) split (different possible variants of option 7) instead, the enhanced CPRI (eCPRI) transport capacities can be reduced compared to CPRI transport by at least a factor of two to three times, and even more when considering the stochastic nature of the traffic at these split points. The numerical simulation of aggregated x-haul traffic in a realistic urban wireless network, taking into account the dynamics of users, their mobility, and download requests, the resulting changes of the radio channel quality, and more, shows how the total midhaul (F1) and fronthaul (7.2x and 7.3) traffic scales with the number of aggregated cells [12]. The total transport capacities are such that, with XGS-PON and 25GS-PON, the business case can still be attractive in terms of number of connected cells per PON (Fig. 9).

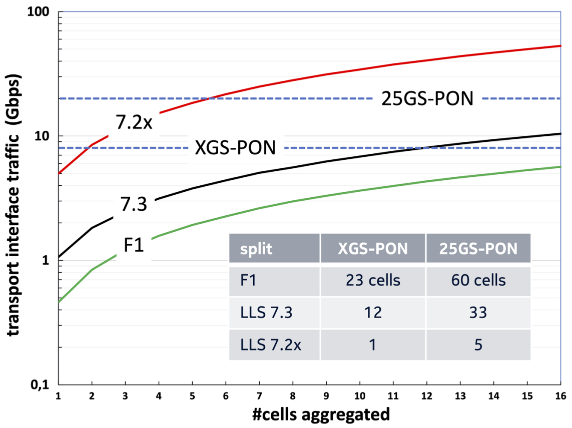

split	XGS-PON	25GS-PON
F1	23 cells	60 cells
LLS 7.3	12	33
LLS 7.2x	1	5

Fig. 9. Aggregated midhaul (F1) and fronthaul (7.2x and 7.3) transport capacities in a simulated urban wireless network, and the respectively supported number of cells per PON, assuming 80% of the nominal PON capacity being usable (see [12] for details).

Regarding the latency requirements on the optical fronthaul network for option 7 splits, they are approximately the same as for the case of option 8 CPRI transport, i.e., in the range of 200 µs round trip (possibly minor differences due the RAN processing functions, e.g., HARQ, being distributed across the nodes in a slightly different way). In contrast to the above case of CPRI transport over TDM-PON, a fixed BW assignment per node would be more inefficient the more dynamic the traffic pattern is at the cell sites. Hence, the cooperative transport interface (CTI) was defined by the O-RAN Alliance [13], which allows for sending a notification from the wireless system to the fixed network transport system (here TDM-PON) about the amount of data from a certain RU to be expected to arrive at the associated ONU for a certain duration of time at a specified future point in time (Fig. 10). On the PON-side, the associated process of pro-active US transport BW assignment is the previously introduced Co-DBA (cf. Fig. 4), which is described in [14]. This information, as well as other configuration data such as, e.g., temporary BW constraints on the PON due to transport of additional services, are exchanged across the bidirectional CTI. Since this notification from the DU to the PON-OLT is sent well in advance, there is only very little buffer waiting time on the ONU, when the data from the RU

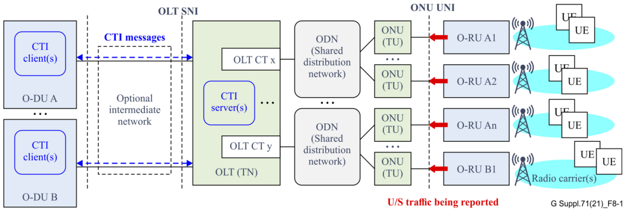

Fig. 10. Converged wireless and PON network architecture including the exchange of messages across the CTI (Fig. 8-1 in [14]).

Table 1. URLLC Services in Several Market Segments and Their Latency and Reliability Requirements [16] (Green Dots Highlight Services with Single Digit Millisecond or Sub-millisecond Latency Requirements)

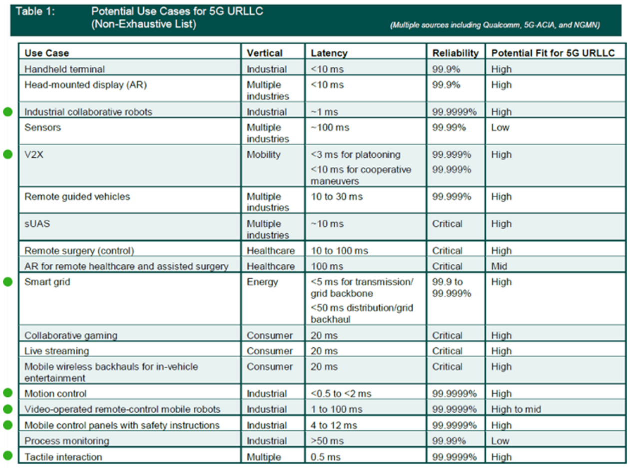

Table 1: Potential Use Cases for 5G URLLC (Non-Exhaustive List)

(Multiple sources including Qualcomm, 5G-ACIA, and NGMN)

	Use Case	Vertical	Latency	Reliability	Potential Fit for 5G URLLC
	Handheld terminal	Industrial	<10 ms	99.9%	High
	Head-mounted display (AR)	Multiple industries	<10 ms	99.9%	High
●	Industrial collaborative robots	Industrial	~1 ms	99.9999%	High
	Sensors	Multiple industries	~100 ms	99.99%	Low
●	V2X	Mobility	<3 ms for platooning <10 ms for cooperative maneuvers	99.999% 99.999%	High
	Remote guided vehicles	Multiple industries	10 to 30 ms	99.999%	High
	sUAS	Multiple industries	~10 ms	Critical	High
	Remote surgery (control)	Healthcare	10 to 100 ms	Critical	High
	AR for remote healthcare and assisted surgery	Healthcare	100 ms	Critical	Mid
●	Smart grid	Energy	<5 ms for transmission/ grid backbone <50 ms distribution/grid backhaul	99.9 to 99.999%	High
	Collaborative gaming	Consumer	20 ms	Critical	High
	Live streaming	Consumer	20 ms	Critical	High
	Mobile wireless backhauls for in-vehicle entertainment	Consumer	20 ms	Critical	High
●	Motion control	Industrial	<0.5 to <2 ms	99.9999%	High
●	Video-operated remote-control mobile robots	Industrial	1 to 100 ms	99.9999%	High to mid
●	Mobile control panels with safety instructions	Industrial	4 to 12 ms	99.9999%	High
	Process monitoring	Industrial	>50 ms	99.99%	Low
●	Tactile interaction	Multiple	0.5 ms	99.9999%	High

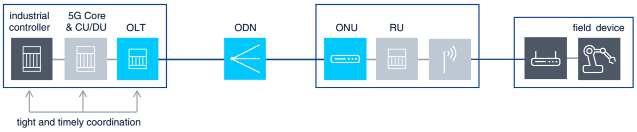

Fig. 11. Tightly coordinated resource assignment across the application, mobile, and fixed network (TDM-PON) transport, here for an industrial use case.

arrive. In principle, it can be as short as the burst period, which is 31.25 µs used in the CPRI proof of concept described in the previous sub-section.

Here again, a similar case can be considered for future IIoT networks. In contrast to the services discussed for IIoT in the previous sub-section, it is now assumed that future IIoT services and other vertical market services will be much more dynamic, will have higher BW requirements, and will be more interconnected [15]. Table 1 shows a selection of such future services, some of which impose end-to-end latency requirements in the range of single digit milliseconds down to below 1 ms (labeled by green dots). This means that the transport system, be it only fixed line transport or including wireless connections, is supposed to be even faster than that.

In the case of a combined wireless/wireline connection, the stringent transport latency across PON can again be addressed by an interconnection architecture as described above, involving the CTI messaging. In some cases (e.g., industrial collaborative robots or tactile interactions), the end-to-end latency may benefit from a tight and timely coordination of the application, the mobile, and the fixed network, using an extended version of CTI-type messaging for real-time coordination (Fig. 11), in addition to a non-real-time network-wide orchestration. This is a topic of future research.

5. PON ARCHITECTURES SUPPORTING LOCAL COMMUNICATION

So far, the discussion in this paper has focused on the optimization of the transport BW assignment in TDM-PON for accommodating low latency services in a ptmp topology, i.e., between a central point (OLT) and multiple end-points (ONUs). There are, however, also use cases in which the end-points should communicate with each other, in addition to the centralized OLT-ONU communication channel. Instead of letting this local communication go through the central and possibly distant OLT, direct passive optical links over the drop fibers would allow for low latency and high BW channels without burdening the OLT–ONU communication channel.

The PON architectures introduced in this section exploit different approaches for establishing such local communication channels for specific use cases [17]. It must be noted that in all cases considered in the following sub-sections, individual parts of the PON architecture need to be modified to some extent. In all three cases, some passive components are added to the fiber infrastructure. In two of them, also the ONU optics are extended to include a second optical receiver (cf. use case "any-to-any") or an additional wavelength division multiplexing (WDM) receiver array and a WDM laser (cf. use case "any-to-few"). These latter add-ons to conventional ONU optics are for special PON variants that may be either standardized in the future or remain a proprietary solution for purpose built networks.

A. Nested PONs ("One-to-Few")

Consider a wireless network in which the antenna sites are located at a distance from the Edge Cloud (>10 km) that does not allow for low latency fronthaul, but only for backhaul or midhaul. In highly densified wireless networks, the macro cells will be supported by several nearby small cells. This arrangement will benefit from close collaboration between macro cells

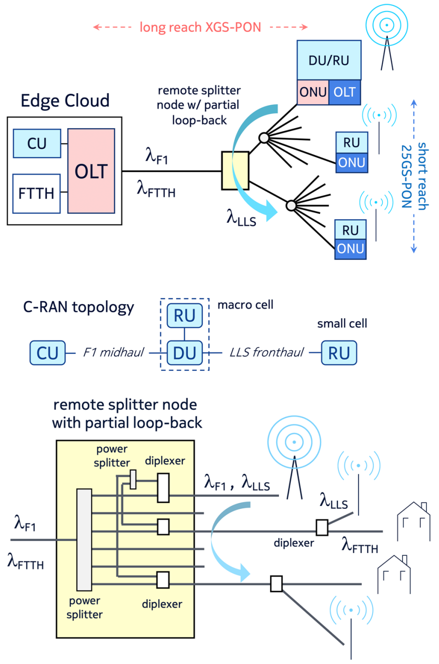

Fig. 12. Nested PON architecture for latency tolerant F1 midhaul over long reach XGS-PON as well as for low latency LLS fronthaul over local short reach 25GS-PON for interconnecting several end-points (top); internal architecture of the remote node (bottom) [18].

and small cells. In addition, the small cell hardware should be reduced as much as possible. In the architecture shown in Fig. 12, the long reach XGS-PON link from the Edge Cloud to the wireless serving area carries the F1 midhaul data for both the macro cell and the associated small cells [18]. At the macro cell site, the DU processes the data for the macro cell RU as well as the data for the small cell RUs. The LLS fronthaul links to the small cell RUs are established by using an additional local 25GS-PON across the ODN drop section together with a slightly modified splitter node (yellow) for selectively looping back the 25GS-PON signals between the macro cell and the small cells. The internal structure of the splitter node comprises an additional low split ratio power splitter for connecting the macro cell with the small cells and diplexers for separating midhaul from fronthaul channels and fronthaul from FTTH channels.

The splitter node can also be implemented in different ways: if the drop fibers to the small cells are not shared with FTTH services, then the diplexers on the drop fibers are not needed inside the node, nor in the field. In an alternative setup, the loopback could be accomplished by using a wavelength selective mirror at the input to the main splitter or at a second input port of the splitter. However, this would introduce unnecessarily high losses of, e.g., >35 dB for a 1:32 main splitter due to the 25GS-PON signals passing through the splitter twice. In contrast, a solution as in Fig. 12 introduces <10 dB for an additional 1:4 splitter plus diplexers, which allows for using low power optics for the 25GS-PON system.

B. Local Grid of Groups of Few Interconnected Nodes ("Any-to-Few")

In a generalized architecture, each node in the ODN drop section is connected to a few nearby neighbors as well to the central OLT (Fig. 13). As an example, a practical use case would be an architecture supporting uplink joint reception in wireless networks with DU/RU on site, in which each cell processes the data of its own antennas and of those in neighboring sites.

In this architecture, the local communication between ONUs is accomplished in a wavelength range different from standardized PON wavelengths. Each ONU continuously broadcasts its data to its connected neighbors on a certain wavelength, e.g., in the E-band, while receiving their respective broadcast data on other, mutually different wavelengths in the same range. Thus, each node has one laser and several receivers in addition to the TDM-PON transceiver used for the connection to the central OLT. The number of required wavelengths for the inter-ONU communication in a given network depends on the detailed topology. The trivial, but inefficient solution for any topology would be "as many wavelengths as nodes," but in certain cases (e.g., small square grids) they can be as few as five wavelengths (= maximum number of neighbors per node + 1). A general approach for finding the most efficient solution with the smallest possible number of wavelengths is for further research.

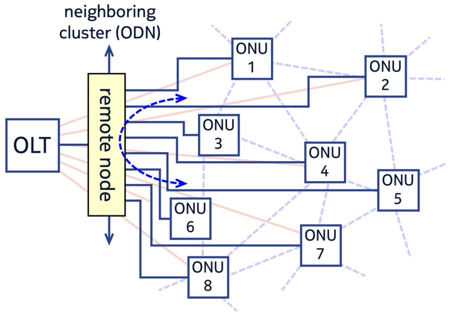

Fig. 13. Generalized ODN architecture (solid blue lines) realizing a conventional ptmp topology (faint solid red lines) and additional logical connections between neighboring ONUs (faint dashed blue lines) by using the ODN drop fibers and a suitably modified remote node (yellow) ([17] copyright Springer).

The interconnection pattern for local traffic is established by adding to the remote node the grid of interconnected power splitters shown in Fig. 14 (top). Each ONU is connected to the common port of its associated 1:4 splitter, while the splitter arms of neighboring ONUs are connected to each other. Unused arms at the edges of the structure can be used for further extending the grid. The interconnection pattern as well as the extension to neighboring ODNs can be flexibly configured by adding optical switching elements inside the white box ([17], not shown here).

The connection to the central OLT is then accomplished by adding the main splitter and fiber links as shown in the middle of Fig. 14. Together, this architecture implements a square grid for local communication within overlapping groups of ONUs, overlaid to the usual ptmp PON topology (bottom of Fig. 14). Connecting the extension ports with their corresponding peers on the opposite side of the white box will yield a ring or a torus topology (for the sake of simplicity, the extension ports on the left and on the right side of the white box have been omitted

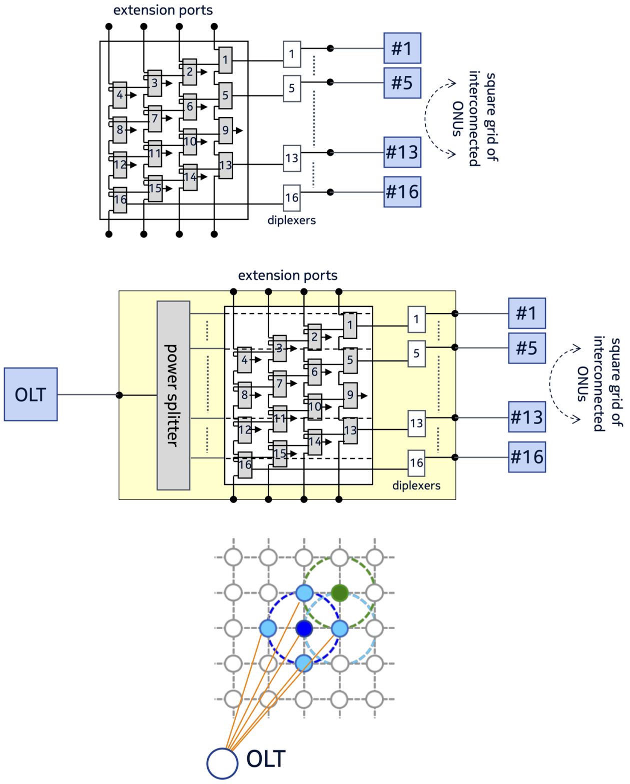

Fig. 14. Grid of interconnected 1:4 power splitters (top) is added to the main PON splitter to form a remote node (middle) that enables inter-ONU communication in overlapping groups of ONUs on a square grid topology (bottom) ([17] copyright Springer).

in the figure). Using splitters with higher split factors in the white box will allow for implementing higher dimensional hyper-cubes and hyper-toruses.

C. Grid for Local Traffic Prevailing over Centralized Traffic ("Any-to-Any")

The third architecture in this section is made for use cases in which local traffic prevails over centralized traffic, such as encountered in intra-data-center networks (east–west versus north–south traffic). To describe the principle of operation, we consider the intra-rack communication between servers in a data center pod. The case of intra-pod communication between racks can be based on a similar concept.

Here, we focus on TDM-PON solutions, so the architecture described in the following will not be sufficient for hyperscale data centers, but it would provide a cost efficient solution for small data centers in future 5G and 6G RAN architectures, far edge clouds, enterprise networks, and more.

The intra-pod architecture in Fig. 15(a) consists of N aggregation switches, N servers per rack, and N top-of-rack switches (ToR) that are here now replaced with N passive remote nodes. The N south-bound ports of the aggregation switches are time/wavelength division multiplexing PON (TWDM-PON) OLT ports, and each server has a TWDM-PON ONU port. Each aggregation switch uses the same wavelength channel pair (DS and US) on all its ports, while different switches use different pairs. The internal architecture of the N × N remote nodes [yellow box in Figs. 15(b)–15(f)] comprises a 2N × 2N star coupler. Each of the N input and output ports of a remote node uses a pair of ports of the internal star coupler (dashed ellipses in Fig. 15).

The sequence of pictures in Fig. 15 shows the working principle. The US channel of a PON is used not only for OLT–ONU communication, but also for the inter-ONU (i.e., inter-server) communication within a rack. The bursts of that US channel are assigned to OLT–ONU communication or to ONU–ONU communication by the TC layer of the OLT port on the respective aggregation switch. The ONUs must include a second receiver that can detect the bursts of the looped back US channel.

Each aggregation switch is connected to each rack in the pod [Fig. 15(a)]. TWDM-PON ONUs are supposed to be wavelength tunable to flexibly join one or the other group of interworking ONUs (servers) in the rack. The control and management plane for this architecture are for further study.

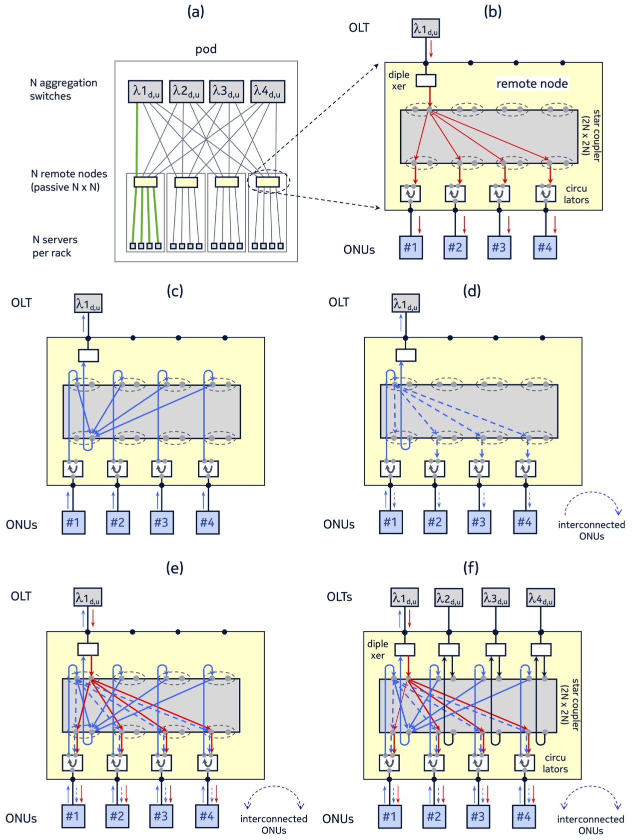

Fig. 15. (a) Intra-pod architecture and optical links for one OLT channel pair (green) through a remote node (yellow box), for ease of reading graphically decomposed into (b) only downstream to all ONUs, (c) only upstream from all ONUs to the OLT, (d) only looped back upstream from ONU #1 to all other ONUs, (e) all three combined, and (f) connecting the remaining OLT ports on aggregation switches #2, #3, and #4 (not showing detailed links inside the 2N × 2N star coupler for those connections) ([17] copyright Springer).

6. CONCLUSION

TDM-PONs are evolving from simple broadband access networks providing only best effort services for residential users and small businesses to a versatile low cost solution also for professional services:

- x-haul transport in wireless networks requiring high capacity, low latency connections, and tight synchronization,
- industrial and vertical markets requiring deterministic networks,
- intra-data-center (small DC) networks requiring low latency intra-rack or intra-pod communication.

This evolution is empowered by steadily increasing system capacities, soon reaching 50G and 100G channel line rates, as well as by flexible multiplexing configurations, supporting statistically multiplexed as well as dedicated fixed bandwidth optical channels in the time domain. By fully exploiting the potential of the standardized TC layer, real-time and deterministic resource assignments can be leveraged for supporting the above services using cost efficient TDM- or TWDM-PON technologies. A major advancement in PON architectures is expected to evolve from the flexibilization of current ptmp communication topologies towards more complex topologies by adding means for enabling intra-PON communication based on local optical channels within the ODN drop section.

REFERENCES

1. "50G PON is the next major step for access networks," Omdia report (2021).
2. "Gigabit-capable passive optical networks (GPON)," ITU-T Recommendation G.984 (2003–2021).
3. "10-Gigabit-capable symmetric passive optical network (XGS-PON)," ITU-T Recommendation G.9807.1 Amd2 (2020).

4. "Higher speed passive optical networks," ITU-T Recommendation G. 9804 series (2021–2022).
5. "40-Gigabit-capable passive optical networks (NG-PON2): Transmission convergence layer specification," ITU-T Recommendation G.989.3 Amd3 (2020).
6. "Multi-source Agreement, 25GS-PON specification v2.0," 2021, https://www.25gspon-msa.org/wp-content/uploads/2021/09/25GS-PON-Specification-V2.0.pdf.
7. "Control, user and synchronization plane specification," ORAN-WG4.CUS.0-v05.00 (O-RAN Alliance, 2021).
8. "F1 general aspects and principles," 3GPP Technical Specification TS 38.470 V16.1.0 (March 2020).
9. R. Bonk, R. Borkowski, M. Straub, H. Schmuck, and Th. Pfeiffer, "Demonstration of ONU activation for in-service TDM-PON allowing uninterrupted low-latency transport links," in *Optical Fiber Communication Conference (OFC)* (2019), paper W3J.4.
10. S. Bidkar, J. Galaro, and Th. Pfeiffer, "First demonstration of an ultra-low-latency fronthaul transport over a commercial TDM-PON platform," in *Optical Fiber Communication Conference (OFC)* (2018), paper Tu2K.3.
11. S. Dietrich, G. May, O. Wetter, H. Heeren, and G. Fohler, "Performance indicators and use case analysis for wireless networks in factory automation," in *IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, Limassol, Cyprus, 2017.
12. S. Bidkar, P. Dom, R. Bonk, and Th. Pfeiffer, "Mobile Xhaul traffic modelling for high-speed TDM-PON," in *European Conference on Optical Communication (ECOC)*, Brussels, Belgium, 2020, paper We2J-5.
13. "Cooperative transport interface, transport control plane specification," ORAN-WG4.CTI-TCP.0-v01.00 (O-RAN Alliance, 2020).
14. "OLT capabilities for supporting CO DBA," ITU-T G.sup71 (2021).
15. "OPC UA: Industrial interoperability standard from sensor to cloud (and back)," OPC Foundation Webinar, December 1, 2020.
16. "5G ultra low latency use cases and market activities," ABIresearch Report AN-5067 (2019).
17. Th. Pfeiffer, "PON architecture enhancements," in *Springer Handbook of Optical Networks* (Springer, 2020), Chap. 29.
18. T. Sizer, D. Samardzija, H. Viswanathan, S. T. Le, S. Bidkar, P. Dom, E. Harstead, and Th. Pfeiffer, "Integrated solutions for deployment of 6G mobile networks," J. Lightwave Technol. (to be published).

Thomas Pfeiffer is with Nokia Bell Labs, working as Principal Architect for optical communication systems and network architectures in access, focusing in particular on converged network solutions for mobile and vertical markets. He has been involved in research, standardization, and product development of optical access systems for more than a decade. Dr. Pfeiffer received his Ph.D. degree while working on ultrafast optoelectronics at the Max-Planck-Institute for Solid State Physics in Stuttgart (Germany). He is a Bell Labs Fellow and a Fellow of Optica.

Pascal Dom is a Senior Product Manager at Nokia with more than 25 years telecommunications experience. He is today a contributor to the passive optical networks (PON) specifications in the area of mobile transport (Xhaul). He holds a master's degree in computer sciences from VUB Brussels Free University (Belgium) and an executive master's degree in business administration (eMBA) from the University of Antwerp (Belgium). He was awarded Distinguished Member of Technical Staff of Nokia as part of the Class of 2021.

Sarvesh Bidkar is a Researcher in the Fixed Networks Research department at Nokia Bell Labs in Stuttgart, Germany. Prior to that, he was a Senior Research Associate at High Performance Networks Group at the University of Bristol from 2014. He received the M.Tech.+Ph.D. degree in computer science and engineering from the Indian Institute of Technology Bombay, India, in 2014. His research interests include resource allocation in multi-domain, multi-technology converged networks, optical network design and management, as well as programmable and low-latency switching and routing. His recent work focuses on optical transport architectures for next-gen radio access networks using low-latency techniques on TDM-PON technology.

François Fredricx is a senior telecom engineer in the Fixed Networks R&D division at Nokia. His focus is on defining and exploiting innovative features in access technologies and platforms for FTTx, from an end-to-end system point of view. Over the course of his career he has contributed to several technical standardization fora (MEF, ITU-T, Broadband Forum, O-RAN Alliance). He has also coordinated several multi-partner EU-funded and local funded projects. Currently he is involved in the standardization and development of PON technologies, especially for the support of Mobile xhaul. He has authored several articles and patents in the field of access technologies.

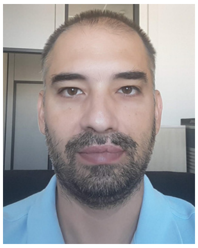

Konstantinos (Kostas) Christodoulopoulos is a researcher at Nokia Bell Labs in Stuttgart, Germany. He was an Adjunct Assistant Professor in the Computer Engineering and Informatics Department of the University of Patras, a senior researcher at the School of Electrical and Computer Engineering of the National Technical University of Athens, Computer Technology Institute, Trinity College Dublin, and a contractor for IBM Dublin. He has authored /co-authored over 100 peer-reviewed journal and conference papers and has over 3000 citations. He has served as an Associate Editor for the *Journal of Optical Communications and Networking* and in various conference committees. His research interests include optical network architectures, static and dynamic optimization, low latency, and deterministic networking.

René Bonk received the Diploma degree in physics from the Technical University of Braunschweig, Germany, and the Dr.-Ing. (Ph.D.) degree in electrical engineering from the Karlsruhe Institute of Technology (KIT), Germany, for work on all-optical signal processing and semiconductor optical amplifiers. In 2012 he joined Alcatel-Lucent, Bell Laboratories, where he focused on next-generation optical metro-access networks and physical layer aspects. He is currently a Department Head of the Converged Access System Research group, Nokia, Bell Laboratories, Stuttgart, Germany. He leads a team that works on PMD and TC-layer aspects of the next-generation passive optical network solutions and new use cases related to convergence. Dr. Bonk has authored and co- authored over 100 peer-reviewed journal and conference papers and a book chapter on "Linear semiconductor optical amplifiers," contributes to IEEE and ITU-T standardization (G.9804.3 co-editor) activities, and has numerous patents/patent applications in the field.